

“Express Mail” Mailing Label No. EL960828023US

PATENT APPLICATION  
ATTORNEY DOCKET NO. SUN-P8774-EKL

5

10       **METHOD AND APPARATUS FOR USING**  
**INTERVAL TECHNIQUES TO SOLVE A**  
**MULTI-OBJECTIVE OPTIMIZATION**  
**PROBLEM**

15

**Inventor:** Gregory R. Ruetsch

**Related Application**

20       [0001] This application hereby claims priority under 35 U.S.C. §119 to  
U.S. Provisional Patent Application No. 60/420,883, filed on 23 October 2002,  
entitled “Interval Algorithm for Multi-Objective Optimization,” by inventor  
Gregory R. Ruetsch.

**BACKGROUND**

25

**Field of the Invention**

[0002] The present invention relates to techniques for solving  
optimization problems within computer systems. More specifically, the present

invention relates to a method and an apparatus that uses interval techniques to solve a multi-objective optimization problem.

### **Related Art**

5           **[0003]** In many real-world optimization problems, there are often several objectives that one would like to optimize. In such cases, rarely do the optima of each objective coincide. Hence, one is left with tradeoffs between the individual objectives that must be incorporated in the solution. The goal of multi-objective optimization is then to determine a set of points that describe the optimal trade-  
10       offs between objectives.

**[0004]** An example of multi-objective optimization is that of minimizing price and maximizing performance. A common metric associated with these two objectives is the price-performance ratio. This number, however, oversimplifies the optimization problem. This oversimplification lies in the implicit assumption  
15       that a unit of price is equal to a unit of performance, which in general is not the case. For example, a college freshman and a national lab would most likely weight these two objectives differently when purchasing a computer. A more insightful approach to the problem is depicted in FIG. 1, which shows the price and runtime (inverse performance) of a set of points representing (fictitious)  
20       computers. Expressed in this manner, one would like to minimize each objective.

**[0005]** Unfortunately, minimizing both objectives simultaneously is generally not possible, and is not the case in this example. The best one can do is come up with an optimal trade-off, where it is not possible to improve one objective without degrading the other. The filled circles in FIG. 1 are points  
25       which exhibit this optimal trade-off, and comprise what is termed the "Pareto front." Obtaining the Pareto front is one of the goals of multi-objective optimization, both in identifying optimal points and in providing a sensitivity

analysis. This sensitivity analysis considers the shape of the Pareto front. Convex portions of the Pareto front (points which lie on the convex hull of all points) indicate an ability to satisfy all objectives relative to regions of the Pareto front which are nonconvex (the “indented” region of Pareto front in FIG. 1).

5           **[0006]** In addition to determining the Pareto front, another goal of multi-objective optimization is to identify values of the objective functions' independent variables which are mapped to points on the Pareto front. In our example, processor speed, memory, system architecture, etc., are all independent variables (perhaps constrained) of both price and performance objective functions. One  
10 would like to know which sets of these independent variables map to points on the Pareto front. Such sets are termed “Pareto optimal sets.”

**[0007]** While the application of intervals to solving the multi-objective optimization problem is new, there are a variety of methods for obtaining the Pareto front for analytical expressions. These methods include the weighted  
15 objective or aggregate function method, evolutionary algorithms (see “Zitzler, E., Deb, D., and L. Thiele (2000) “Comparison of Multiobjective Evolutionary Algorithms: Empirical Results,” *Evolutionary Computing* 8(2): 173-195), and adjoint-gradient methods (see Nemec, M. and D. W. Zingg (2001) “Towards Efficient Aerodynamic Shape Optimization Based on the Navier-Stokes  
20 Equations,” *15<sup>th</sup> AIAA Computational Fluid Dynamics Conference*, AIAA 2001-2532). Each of these methods has its benefits and detractions. The aggregated function method turns the problem into a single objective optimization by weighting each objective. For example, if we have  $n$  objectives,  $\{f_1, \dots, f_n\}$ , we wish to simultaneously minimize, then one can minimize the fitness function

25

$$F = \sum_{i=1}^n w_i f_i(\mathbf{x})$$

for all combinations of  $w_i$  such that  $0 \leq w_i \leq 1$  and

$$\sum_{i=1}^n w_i = 1 .$$

5

[0008] Although simple in concept, there are several disadvantages to this method. The first is that uniform variations in the weightings space do not evenly map out the Pareto front. A second, and more substantial, disadvantage is that this method cannot capture nonconvex regions of the Pareto front. A third  
10 disadvantage is that the method does not guarantee finding all solutions. (Note that this disadvantage is to common to all existing solution methods.)

[0009] Hence, what is needed is a method and an apparatus for determining a Pareto front for a multi-objective optimization problem without the above-described problems.

15

## SUMMARY

[0010] One embodiment of the present invention provides a system that uses interval techniques to solve a multi-objective optimization problem. During operation, the system receives a representation of multiple objective functions  
20  $(f_1, \dots, f_n)$  at a computer system, wherein  $(f_1, \dots, f_n)$  are scalar functions of a vector  $\mathbf{x} = (x_1, \dots, x_n)$ . The system also receives a representation of a domain of interest for the multiple objective functions. Next, the system performs an interval optimization process to compute guaranteed bounds on a Pareto front for the objective functions  $(f_1, \dots, f_n)$ , wherein for each point on the Pareto front, an  
25 improvement in one objective function cannot be made without adversely affecting at least one other objective function. While performing the interval

optimization process, the system applies a direct-comparison technique between subdomains of the domain of interest to eliminate subdomains that are certainly dominated by other subdomains.

5 [0011] In a variation on this embodiment, performing the interval optimization process additionally involves applying a gradient technique to eliminate subdomains that do not contain local Pareto optima. A subdomain  $[\mathbf{x}]_i$  is eliminated by the gradient technique if an intersection of certainly negative gradient regions  $\mathbf{C}_j$  for each objective function  $f_j$  is non-empty,  $\bigcap_{j=1}^n \mathbf{C}_j([\mathbf{x}]_i) \neq \emptyset$ , wherein the certainly negative gradient region  $\mathbf{C}_j$  for objective function  $f_j$  is the  
10 intersection of  $\underline{\mathbf{N}}_j([\mathbf{x}]_i)$  (the negative gradient region associated with the minimum angle  $\underline{\theta}_j$  of the gradient of  $f_j$  over the subdomain  $[\mathbf{x}]_i$ ) and  $\overline{\mathbf{N}}_j([\mathbf{x}]_i)$  (the negative gradient region associated with the maximum angle  $\overline{\theta}_j$  of the gradient of  $f_j$  over the subdomain  $[\mathbf{x}]_i$ ).

[0012] In a further variation, the system performs an iterative process that,  
15 (1) bisects remaining subdomains that have not been eliminated by the gradient technique, and (2) applies the gradient technique to eliminate bisected subdomains that do not contain a local Pareto optimum.

[0013] In a further variation, bisecting a subdomain involves bisecting the subdomain in the direction that has the largest width of partial derivatives of all  
20 objective functions  $(f_1, \dots, f_n)$  over the subdomain.

[0014] In a further variation, the direct-comparison technique is applied once for every  $n$  iterations of the gradient technique.

[0015] In a further variation, the iterations continue until either a predetermined maximum number of iterations are performed, or the largest area of  
25 any subdomain is below a predetermined value.

[0016] In a variation on this embodiment, a subdomain  $U$  certainly dominates a subdomain  $V$  if every point  $u \in U$  dominates every point  $v \in V$ , wherein a point  $u$  dominates a point  $v$  under minimization if,  $u_i \leq v_i$ ,  $i = 1, \dots, n$ , and  $u_i < v_i$  for some  $i \in \{1, \dots, n\}$ .

5

## BRIEF DESCRIPTION OF THE FIGURES

[0017] FIG. 1 presents a graph of price versus performance for computer systems.

[0018] FIG. 2 depicts a Pareto front.

10 [0019] FIG. 3 illustrates a Pareto optimal set.

[0020] FIG. 4A illustrates a non-positive gradient region for a first objective function in accordance with an embodiment of the present invention.

[0021] FIG. 4B illustrates a non-positive gradient region for a second objective function in accordance with an embodiment of the present invention.

15 [0022] FIG. 4C illustrates overlapping non-positive gradient regions in accordance with an embodiment of the present invention.

[0023] FIG. 5A illustrates a first non-positive gradient region for a point which is Pareto optimal in accordance with an embodiment of the present invention.

20 [0024] FIG. 5B illustrates a second non-positive gradient region for the point which is Pareto optimal in accordance with an embodiment of the present invention.

[0025] FIG. 6A illustrates a non-positive gradient region associated with a minimum gradient angle for a subdomain in accordance with an embodiment of the present invention.

25

[0026] FIG. 6B illustrates a non-positive gradient region associated with a maximum gradient angle for the subdomain in accordance with an embodiment of the present invention.

5 [0027] FIG. 6C illustrates a *certain* non-positive gradient region for the subdomain in accordance with an embodiment of the present invention.

[0028] FIG. 7 presents a flow chart of an interval optimization process for a multi-objective function in accordance with an embodiment of the present invention.

10 [0029] FIG. 8 presents a more-detailed flow chart for the interval optimization process in accordance with an embodiment of the present invention.

## DETAILED DESCRIPTION

[0030] The following description is presented to enable any person skilled in the art to make and use the invention, and is provided in the context of a particular application and its requirements. Various modifications to the disclosed  
15 embodiments will be readily apparent to those skilled in the art, and the general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the present invention. Thus, the present invention is not intended to be limited to the embodiments shown, but is  
20 to be accorded the widest scope consistent with the principles and features disclosed herein.

[0031] The data structures and code described in this detailed description are typically stored on a computer readable storage medium, which may be any device or medium that can store code and/or data for use by a computer system.  
25 This includes, but is not limited to, magnetic and optical storage devices such as disk drives, magnetic tape, CDs (compact discs) and DVDs (digital versatile discs or digital video discs), and computer instruction signals embodied in a

transmission medium (with or without a carrier wave upon which the signals are modulated). For example, the transmission medium may include a communications network, such as the Internet.

## 5 **Interval Computations**

[0032] Note that interval computation provides an automatic means for bounding roundoff error. Rather than representing a number by a single floating point value, which may or may not exist in machine representation, an interval represents a number by machine representable lower and upper bounds which  
10 contain the desired value. Interval computation uses directed or outward rounding such that the resulting interval always contains the mathematical result. In addition to bounding roundoff error, interval computation provides a rich environment for performing set theoretic operations. One of the most powerful aspects of interval computation, however, is that in performing a single function  
15 evaluation over an interval, one obtains rigorous bounds of that function *over the entire interval*. Depending on the value of these bounds, one can prove or disprove certain conjectures. We leverage this feature in multi-objective optimization to eliminate subdomains in which we can prove no point is Pareto optimal. In doing so, we capture all Pareto optimal points. In essence, intervals  
20 allow one to develop computer techniques, which deal with continua rather than discrete points, and as a result one can perform proofs via computers.

[0033] In the next section we utilize the objective function derivatives to formulate a differential definition for local Pareto optima, which will become the basis for the interval-differential formulation. We later present two  
25 complimentary techniques for solving the multi-objective problem: (1) a direct-comparison technique, which can be used to solve the global problem, but suffers from inefficiency, and (2) an interval version of the differential formulation,



which efficiently solves for local optima. The combination of these two interval techniques results in a composite technique, which is guaranteed to find all global Pareto optimal sets and their projection on the Pareto front.

## 5 **Multi-Objective Concepts and Definitions**

[0034] In this section, we briefly formalize concepts and define terms used in multi-objective optimization. We begin by presenting an analytical example, and once the multi-objective optimization concepts have been established through this example, we generalize the concepts.

- 10 [0035] Consider a case with two objective functions which represent the distance of a point  $(x_1, x_2)$  from points  $(-1/2, 0)$  and  $(1/2, 0)$ ,

$$f_1(x_1, x_2) = \sqrt{\left(x_1 + \frac{1}{2}\right)^2 + x_2^2} \quad f_2(x_1, x_2) = \sqrt{\left(x_1 - \frac{1}{2}\right)^2 + x_2^2} \quad (1)$$

- 15 which we want to minimize in a two-dimensional domain

$$(x_1, x_2) \in ([-1, 1], [-1/2, 1/2]).$$

- [0036] The individual minima of  $f_1$  and  $f_2$  are the points  $(-1/2, 0)$  and  $(1/2, 0)$ , respectively, so the objectives are minimized at different points in the domain, as is generally the case in multi-objective optimization. What we hope to achieve for such a set of objective functions is an optimal trade-off in objectives. This optimal trade-off can be determined by projecting all the points in this domain into objective space, as depicted in FIG. 2.

- 25 [0037] In FIG. 2 the projection of points  $(x_1, x_2) \in ([-1, 1], [-1/2, 1/2])$  into objective space  $(f_1, f_2)$  is indicated by the crosshatched region. (Some points are projected above and to the right of the crosshatched region). The points which lie on the bold line comprise the Pareto front for this pair of functions. Points on the

Pareto front are nondominated, in that no point exists that can improve one objective without degrading the other.) When projected to objective space, the points in this domain fill the region above the diagonal line. The diagonal line is called the “Pareto front” and represents the optimal trade-off in that there can be  
5 no improvement towards minimizing one objective which doesn't adversely affect the other objective. (Since the minimum of each objective is zero, this condition can be interpreted geometrically by requiring the rectangle whose diagonal connects the origin and the point in question to be empty.) On the other hand, all points which lie above the Pareto front can simultaneously improve both  
10 objectives by moving towards the Pareto front.

[0038] In addition to determining the Pareto front, the corresponding points in the domain  $(x_1, x_2)$  which project onto the Pareto front are desired. Such a point is called “Pareto optimal,” and the collection of all such points the “Pareto optimal set.” For our example, the Pareto optimal set includes all points in the line  
15 segment connecting the two individual minima, depicted in FIG. 3. FIG. 3 illustrates contours (dashed lines) and the Pareto optimal set (solid line) for functions  $f_1$  and  $f_2$  in Equation (1). The value of these functions are the distances from points  $(-1/2, 0)$  and  $(+1/2, 0)$ , respectively, which result in the circular contours. The Pareto optimal set for this example comprises the points on the line  
20 connecting the two individual minima, and projects into the Pareto front in objective space. Note that any point in the domain that is not in the optimal Pareto set can improve both objectives by moving along a path towards the line which comprises the Pareto optimal set.

[0039] With the concepts of Pareto front and Pareto optimal set for the  
25 above analytical expressions established, we now turn our attention to generalizing these concepts to a case with  $n$  objective functions in a  $D$ -dimensional domain. In order to generalize the Pareto front to  $n$  dimensions, it is

useful to establish the concept of dominance. If  $\mathbf{u}$  and  $\mathbf{v}$  are  $n$ -dimensional vectors, then under minimization  $\mathbf{u}$  is said to dominate  $\mathbf{v}$ , denoted as  $\mathbf{u} < \mathbf{v}$ , if:

$$\begin{aligned} &u_i \leq v_i, i = 1, \dots, n; \text{ and} \\ &u_i < v_i \text{ for some } i \in \{1, \dots, n\}. \end{aligned}$$

[0040] Applying dominance to the Pareto concepts, we say the Pareto front contains points which are *nondominated* solutions: solutions where an improvement in one objective cannot be made without adversely affecting at least one other objective. Generally, if we wish to minimize  $n$  objectives  $\{f_1, \dots, f_n\}$  within a particular domain  $\Omega$  of a  $D$ -dimensional space  $\mathbf{x} = (x_1, \dots, x_D)$ , and we define an objective vector  $f(\mathbf{x}) = (f_1(\mathbf{x}), \dots, f_n(\mathbf{x}))$ , then a point  $\mathbf{x} \in \Omega$  is Pareto optimal if:

$$\mathbf{f}(\mathbf{x}') \nless \mathbf{f}(\mathbf{x}) \quad \forall \mathbf{x}' \in \Omega \quad (2)$$

[0041] The definition of Pareto optima in Equation (2) utilizes the function evaluations. This point definition becomes the basis for the interval direct-comparison method in a following section. In the next section we utilize the objective function derivatives to formulate a differential definition for local Pareto optima, which will become the basis for the interval-differential formulation.

### **Differential Formulation**

[0042] If the objective functions  $\{f_1(\mathbf{x}), \dots, f_n(\mathbf{x})\}$  are continuous and differentiable, then it is helpful to leverage the gradient information of these functions in determining *local* Pareto optima. Just as finding local extrema in single-objective optimization does not guarantee that these points are global extrema, finding local Pareto optima does not imply that these points are

contained in the Pareto optimal set. However, finding points that are local Pareto optima can assist in obtaining the Pareto optimal set. This is especially true in the case of interval methods, where consistency checks can be combined with capturing local optima to determine global optima.

- 5           **[0043]** Local Pareto optimality requires that at a point  $\mathbf{x}$ , *no* differential displacement vector  $\delta\mathbf{r}$  exists such that:

$$\mathbf{f}(\mathbf{x} + \delta\mathbf{r}) < \mathbf{f}(\mathbf{x}).$$

- 10    If we denote  $\mathbf{d}(\mathbf{x}, \delta\mathbf{r})$  as the vector of directional derivatives of the objective functions along  $\delta\mathbf{r}$ :

$$d_i(\mathbf{x}, \delta\mathbf{r}) = \nabla f_i|_{\mathbf{x}} \cdot \delta\mathbf{r},$$

- 15    then a Taylor expansion of  $\mathbf{f}(\mathbf{x})$  results in the condition for local Pareto optimality: a point  $\mathbf{x}$  is locally Pareto optimal if there is no differential displacement vector  $\delta\mathbf{r}$  such that:

$$\mathbf{d}(\mathbf{x}, \delta\mathbf{r}) < \mathbf{Z} \tag{3}$$

- 20    where  $\mathbf{Z}$  is the zero vector, i.e.  $Z_i = 0$  for  $i = 1, \dots, n$ .

- [0044]** The condition of Equation (3) can be geometrically interpreted as follows. For each objective there is a  $(D-1)$ -dimensional plane which is normal to the gradient of  $f_i$  at a point  $\mathbf{x}$ . This plane delineates two semi-infinite regions, one where the local (linearized) gradient is positive, and one where it is negative.
- 25    Any differentially small displacement into the negative gradient region results in improving the objective function, and Equation (3) is satisfied when the intersection of the such regions, one for each objective, is empty. If we represent the negative gradient region for  $f_i$  at point  $\mathbf{x}$  by  $N_i(\mathbf{x})$ , then Equation (3) is satisfied
- 30    when:

$$\bigcap_{i=1}^n N_i(\mathbf{x}) \cap \Omega = \emptyset \quad (4)$$

The domain  $\Omega$  appears here in order to account for points which lie on the boundary of  $\Omega$ . In effect, for points on the boundary, this acts like an additional objective function. For interior points, the presence of  $\Omega$  is inconsequential.

[0045] The application of the condition in Equation (4) to the example problem in Equation (1) is demonstrated in FIGs. 4A-4C. In FIGs. 4A and 4B, the objectives  $f_1$  and  $f_2$  are dealt with separately. In each plot, the gradient of the objective at some point is drawn, along with  $N_i \cap \Omega$  (the intersection of the negative gradient region and the domain) by the shaded regions. In FIG. 4C, these two figures are combined, showing a nonempty region corresponding to  $N_1 \cap N_2 \cap \Omega$ . Any differential displacement into this crosshatched region will improve both objectives, and as a result the point is not locally Pareto optimal.

[0046] FIGs. 5A and 5B illustrate negative gradient regions  $N_i$  at a point which is Pareto optimal.  $N_1$  and  $N_2$  are shown independently in the two diagrams. Here the gradients are co-linear but opposite in direction, and as such the intersection  $N_1 \cap N_2$  is empty. Therefore, this point is locally Pareto optimal. Note that the only way an interior point in a two-objective optimization problem can be Pareto optimal is when the gradients are co-linear but opposite.

[0047] In this section, we have established a differential formulation for identifying local Pareto optima. Application of this methodology as it stands is difficult in that one needs to perform this evaluation at every point in domain  $\Omega$ . One can devise many “shotgun” methods using this differential formulation, such as seeding the domain with points and using the gradient information to trajectory points towards the optimal Pareto set. The problem with such methods is that they are not guaranteed to capture all Pareto optima. In the next section, we apply

interval methods to the differential formulation which, along with in interval version of Equation (2), results in a technique which is *guaranteed* to capture all Pareto optima.

## 5 **Interval Concepts**

[0048] Interval methods grew out of the desire to control roundoff error in floating point computations. The idea is that rather than using a single floating point value to represent a number, which would incur an error if the number is not machine representable, a number is represented by upper and lower bounds which are machine representable. Interval arithmetic is performed in a way to guarantee containment. Intervals are described in more detail in U.S. Patent No. 6,629,120, entitled "Method and Apparatus for Performing a Mask-Driven Interval Multiplication Operation," by inventors William G. Walster and Dmitri Chirayev.

[0049] In addition to operators for arithmetic and function evaluation, intervals have a rich set of relational operators. There are *certainly* true operations, for example:

$$\mathbf{X} \overset{c}{<} \mathbf{Y} \Leftrightarrow x < y \quad \forall x \in \mathbf{X} \text{ and } y \in \mathbf{Y}$$

as well as *possibly* true operations:

$$\mathbf{X} \overset{p}{<} \mathbf{Y} \Leftrightarrow \exists x \in \mathbf{X} \text{ and } y \in \mathbf{Y} \mid x < y$$

These same ideas hold for all relational operators.

[0050] Now that some basic concepts of interval arithmetic and computation have been discussed, we apply these concepts to multi-objective optimization, starting with the interval definition of dominance.

### Interval Dominance And The Direct-comparison Method

[0051] We can extend the notions of *certainly* and *possibly* used to define the basic interval relational operators to that of dominance. An interval vector  $U = (U_1(x), \dots, U_n(x))$  certainly dominates an interval vector  $V$  according to the following definition:

$$U \overset{c}{\prec} V \Leftrightarrow u \prec v \quad \forall u \in U \text{ and } v \in V \quad (6)$$

In terms of the infima and suprema of the components of  $U$  and  $V$ , we say that  $U \overset{c}{\prec} V$  if:

$$\begin{aligned} \overline{u_i} &\leq \underline{v_i} \quad i = 1, \dots, n \text{ and} \\ \overline{u_i} &< \underline{v_i} \text{ for some } i \in \{1, \dots, n\}. \end{aligned}$$

Likewise,  $U$  is said to *possibly dominate*  $V$ , denoted as  $U \overset{p}{\prec} V$  if:

$$\underline{u_i} \leq \overline{v_i} \quad i = 1, \dots, n.$$

One could apply the definition of certain dominance directly to obtain an interval technique for determining the Pareto front. One would simply search over a set of subdomains covering the domain  $\Omega$ , and eliminate any subdomains which were certainly dominated. Then the remaining subdomains would be bisected, and the procedure is repeated until some stopping criterion is met.

[0052] While this method will work in that no Pareto optimal points will be eliminated, there are several disadvantages inherent in this method. The main disadvantage is that the number of comparisons to determine certain dominance grows rapidly with number of subdomains. Performance-wise this becomes costly as the Pareto front is refined.

[0053] One aspect of this direct-comparison method that is both advantageous and disadvantageous is that it is a global method. This is disadvantageous with respect to parallelization in a clustered environment--some message passing is needed for comparison between all subdomains within an iteration. The global nature of the direct-comparison method is advantageous in that it results in global Pareto optima. This is important since the interval version of the differential formulation produces all the local Pareto optima, and some consistency check is needed to eliminate those local optima which are not global. The direct-comparison method provides this consistency check. Having developed the direct-comparison method, the only remaining piece in obtaining a complete technique is to extend the differential formulation to intervals.

#### **Interval-differential Formulation**

[0054] In the direct-comparison method above, we used the concept of *certain dominance* through the evaluation of the objectives over intervals to eliminate subdomains, which are certainly dominated, and therefore contain no Pareto optimal points. In the interval-differential formulation, we again look to develop a criterion which enable us to remove subdomains which are certainly dominated, however this criterion is based on the objective function derivatives.

[0055] In the point differential formulation, determining local Pareto optima requires the evaluation of the gradient of objective functions. The magnitude of the gradients is inconsequential to the formulation, it is the direction of the gradients that determined the negative gradient regions,  $N_i$ , and therefore local dominance. In a  $D$ -dimensional domain, the direction of the gradient of each objective function can parameterized by  $D-1$  variables. For example, in the two-dimensional domain of Equation (1), we can define the direction of the gradient of the  $j$ -th objective by the angle  $\theta_j$ . If we extend this to intervals, then the interval



gradient of the  $j$ -th objective function evaluated over the  $i$ -th subdomain,  $[\mathbf{x}]_i$ , is:

$$\mathbf{G}_j([\mathbf{x}]_i) = [\nabla f_j](\mathbf{x})_i$$

5 takes on a range of directions  $[\underline{\theta}_j, \overline{\theta}_j]$ . Associated with this range of directions is a range of semi-infinite negative gradient region, denoted as  $\mathbf{N}_j$ . Formally we have:

$$N_j(\mathbf{x}) \subset \mathbf{N}_j([\mathbf{x}]_i) \quad \forall \mathbf{x} \in [\mathbf{x}]_i$$

10

In the point differential formulation, regions where the intersection of the domain and all  $N_i$  is non-empty indicated that such a point is not Pareto optimal. We now extend this idea to the interval case. To eliminate a subdomain, we need to establish that every point  $\mathbf{x}$  in this subdomain satisfies:

15

$$\bigcap_{j=1}^n N_j(\mathbf{x}) \neq \emptyset. \quad (7)$$

The intersection with the entire domain  $\Omega$  is not included here, as it was in the point formulation--we consider boundary points later in this section. We would like to find an expression analogous to Equation (7) using  $\mathbf{N}_j$ . Suppose we can determine a *certainly negative gradient* region for objective  $f_j$ , denoted by  $\mathbf{C}_j$ , which is defined as the subset of  $\mathbf{N}_j$  that is common to  $N_j(\mathbf{x})$  for all  $\mathbf{x} \in [\mathbf{x}]_i$ :

20

$$\mathbf{C}_j([\mathbf{x}]_i) \subset N_j(\mathbf{x}) \quad \forall \mathbf{x} \in [\mathbf{x}]_i$$

25

With knowledge of  $\mathbf{C}_j$  we can guarantee that Equation (7) holds at all points in the subdomain if:

$$\bigcap_{j=1}^n \mathbf{C}_j([\mathbf{x}]_i) \neq \emptyset.$$

The remaining issue is to determine  $C_j([x]_i)$ . This region is drawn for a single objective,  $f_1$ , in FIGs. 6A-6B. FIG. 6A shows the negative gradient region of  $f_1$  associated with the minimal angle  $\underline{\theta}_1$ , of the gradient over the subdomain  $[x]_i$ .

- 5 FIG. 6B shows the negative gradient region of  $f_1$  associated with the maximum angle  $\overline{\theta}_1$  of the gradient over this subdomain. These negative gradient regions are denoted by  $\underline{N}_1$  and  $\overline{N}_1$ , respectively. FIG. 6C shows the intersection of these two regions, which defines the certainly negative gradient region for objective  $f_1$ . We can define such a region for all objectives:

10

$$C_j = \underline{N}_j([x]_i) \cap \overline{N}_j([x]_i)$$

As a result, we now have a condition which, when satisfied, shows that every point in a subdomain  $[x]_i$  is not locally, and therefore not globally, Pareto optimal:

15

$$\bigcap_{j=1}^n \left( \underline{N}_j([x]_i) \cap \overline{N}_j([x]_i) \right) \neq \emptyset. \quad (8)$$

- [0056] The boundary of  $\Omega$  in the point formulation was accounted for by taking the intersection of the  $N_i$ 's with the domain  $\Omega$ . When operating over a finite area subdomain, there is no extension to account for the boundary in such a simple manner. One needs to account for the boundaries separately. For example, for two objective functions  $f_1$  and  $f_2$  in a domain of  $x \in ([-1,1], [-1/2, 1/2])$ , one needs to solve two problems. The first using  $C_1 \cap C_2$ , or Equation (8), as a criterion for eliminating subdomains in the interior, and the second using
- 20  $C_1 \cap C_2 \cap \Omega$  as a criterion for eliminating subdomains of the semi-degenerate intervals  $(-1, [-1/2, 1/2])$ ,  $(1, [-1/2, 1/2])$ ,  $([-1,1], -1/2)$ , and  $([-1,1], 1/2)$ . An application
- 25

of the direct-comparison method over all subdomains from these various regions merges these separate optimization problems.

[0057] The criterion in Equation (8) for eliminating subdomains which do not contain any local Pareto optima, along with the definition of interval dominance of Equation (6) which be used to eliminate subdomains which are not global Pareto optimal, can be combined in a robust and efficient implementation of a multi-objective optimization code. Equation (8) provides an efficient means for computing local Pareto optima, and Equation (6) can be applied as a consistency check to eliminate any regions which are local but not global optima.

10

### **System Implementation**

[0058] One embodiment of the present invention implements a system that uses the interval-differential formulation, from Equation (8), and the direct-comparison method, from Equation (6). This embodiment includes a single loop, which performs eliminations of subdomains within each “generation” or iteration.

[0059] Referring to the flowchart in FIG. 7, the system first receives a representation of multiple objective functions ( $f_1, \dots, f_n$ ) and a domain of interest (step 702). Next, the system stores this representation in a memory in a computer system (step 704), and performs an interval optimization process to compute guaranteed bounds on the Pareto front for the multiple objective functions (step 706).

[0060] The initial condition, or first iteration, of this loop operates on a single “subdomain” which covers the entire domain of interest. Referring to the flowchart in FIG. 8, every few iterations through the loop a consistency check for global Pareto optima is performed using the direct-comparison method from Equation (6) (step 802). This eliminates any subdomains which are certainly dominated (step 802). The frequency at which the direct-comparison method

occurs is a parameter set by the user.

[0061] Next, the loop determines if the subdomains can be eliminated on grounds of not containing a local Pareto optimum via the interval-differential formulation in Equation (8) (step 804). The certainly negative gradient regions  
5 are represented by interval angles.

[0062] To accommodate problems of more than two independent variables, one simply considers the projections of  $\underline{N}_j([\mathbf{x}]_i)$  and  $\overline{N}_j([\mathbf{x}]_i)$  on all possible two-dimensional coordinate planes. For the case of two independent variables, one finds certainly negative gradient regions on the  $x_1$ - $x_2$  plane. In three  
10 dimensions this is done on three planes, ( $x_1$ - $x_2$ ,  $x_1$ - $x_3$ , and  $x_2$ - $x_3$ ), and in  $D$  dimensions this operation is performed on  $D(D-1)$  such planes. Since the semi-infinite regions  $\underline{N}_j([\mathbf{x}]_i)$  and  $\overline{N}_j([\mathbf{x}]_i)$  are bounded by  $D$ -dimensional planes, the projection of these regions into each two-dimensional subspace is simply done using the appropriate partial derivatives for the subspace to form the gradient, and  
15 then proceeding as if it were a two-dimensional problem.

[0063] The intersection of the certainly negative gradient regions within each coordinate plane is performed using these interval angles, accounting for branch cuts when required. While the number of coordinate planes is  $O(D^2)$  in number of dimensions, one must remember that the operation performed on each  
20 coordinate plane is a simple intersection. The evaluation of the objective gradients, which in general is more time-consuming, is only  $O(D)$ .

[0064] If the intersection of certainly negative gradient regions for all such planes is nonempty, then the subdomain can be removed, meaning they are not advanced into the next iteration. The subdomains that do not satisfy this  
25 condition are bisected (step 806). The bisection is performed in the direction that has the largest width of partial derivatives of all objective functions over the

subdomain. The rationale behind bisecting in this direction is that we want the next iteration of subdomains to have the certain negative gradient regions as large as possible, thus increasing the likelihood that  $C_i \neq \emptyset$ . This is achieved if we bisect a subdomain in the direction where the largest width of the objective functions' partial derivative occurs.

[0065] At the end of the loop the system determines if a termination condition is satisfied (step 808). For example, the termination condition can be satisfied if either a specified number of iterations is performed, or the largest area of any interval box ( $F_1([x]_i, F_2([x]_i)$ ) is below some user specified value. If the termination condition is satisfied, the process is complete. Otherwise, the system returns to step 802 to perform another loop iteration.

[0066] The foregoing descriptions of embodiments of the present invention have been presented for purposes of illustration and description only. They are not intended to be exhaustive or to limit the present invention to the forms disclosed. Accordingly, many modifications and variations will be apparent to practitioners skilled in the art. Additionally, the above disclosure is not intended to limit the present invention. The scope of the present invention is defined by the appended claims.